



CMPSCI 670: Computer Vision


Texture

University of Massachusetts, Amherst
October 6, 2014

Instructor: Subhransu Maji

Administrivia


- Homework 2 ~~is~~ was due today
- Homework 3 posted!
 - implement a “blob detector”
 - due on October 20



Recap: last few lectures



- Convolution
 - Linearity and separability
- Edge detection
 - Find locations where there is high derivatives
 - Canny edge detector - linking weak edges with strong edges
- Corner detection
 - Find locations where intensity changes rapidly in all directions
- Blob detection (scale covariant detector)
 - Convolve with a Laplacian of Gaussian at multiple scales
 - Find maxima over scale and space

Texture




widespread, easy to recognize, but hard to define

Includes: more regular patterns



Includes: more random patterns






Texture-related tasks

- **Shape from texture**

- Estimate surface orientation or shape from image texture

Shape from texture

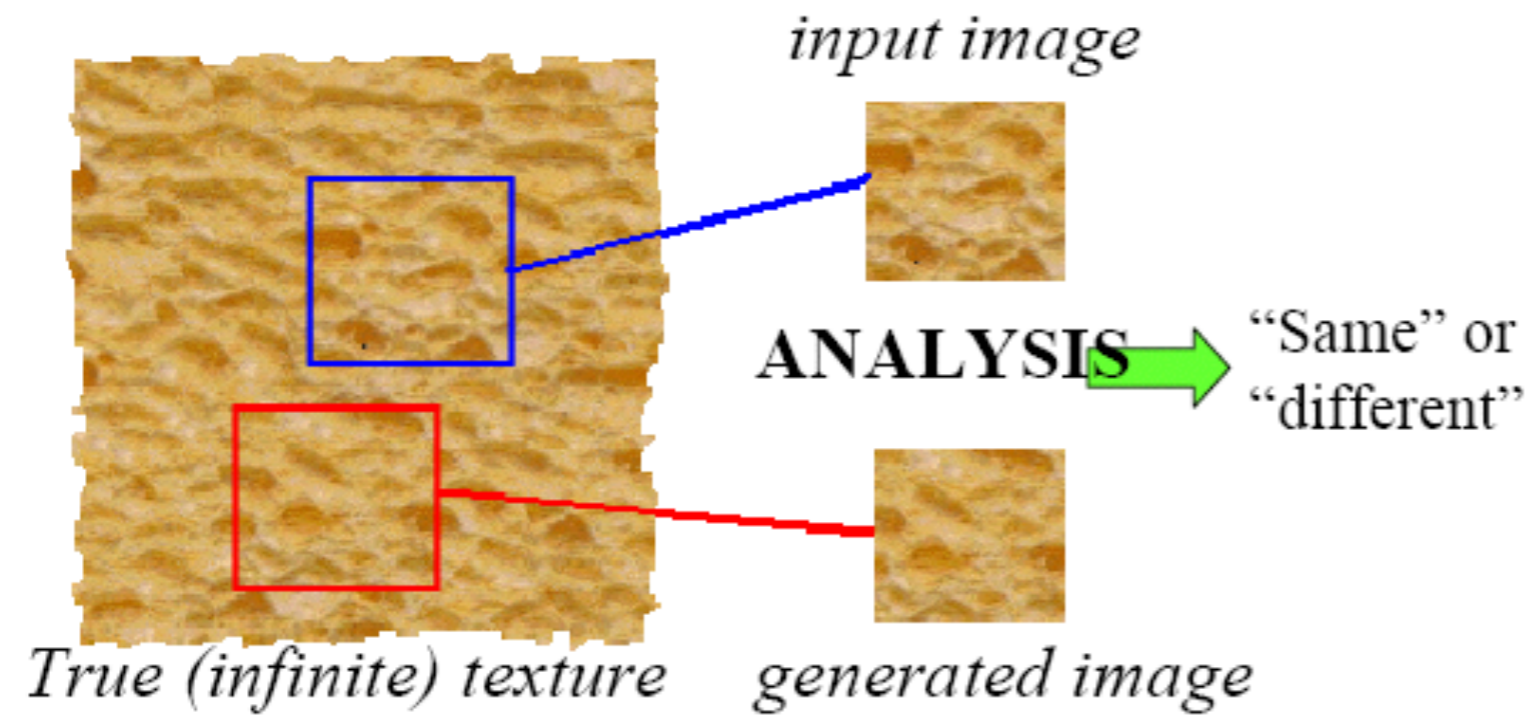
- Use deformation of texture from point to point to estimate surface shape




Texture-related tasks

- **Shape from texture**
 - Estimate surface orientation or shape from image texture
- **Segmentation/classification** from texture cues
 - Analyze, represent texture
 - Group image regions with consistent texture
- **Synthesis**
 - Generate new texture patches/images given some examples


Analysis vs. Synthesis




Why analyze texture?



Texture is indicative of material



.. of object type, especially when shape is not useful



.. of object type, especially when shape is not useful



Why analyze texture?


- Importance to perception:
 - Often indicative of a material's properties, e.g. shiny vs. rough. There is evidence that we can do this using visual cues only (Edelson et al.)
 - Can be important appearance cue, especially if shape is similar across objects
 - Aim to distinguish between occlusion boundaries and texture — good for recognition.
- **Technically:**
 - Representation-wise, we want a feature one step above “building blocks” of corners, blobs and edges.

Psychophysics of texture

- Some textures distinguishable with **pre-attentive** perception – without scrutiny, eye movements [Julesz 1975]



Same or different?

L L L L L L L L L L L L
L L L L L L L L L L L L
L L L L L L L L L L L L
L L L L L L L L L L L L
L L L L L L L L L L L L
L L L L L L L L L L L L
L L L L L L L L L L L L



Textons


“local” unit of texture






Texture representation

- Textures are made up of repeated local patterns, so:
 - Find the patterns
 - Use filters that look like patterns
 - e.g. spots, edges, bars
 - Consider magnitude of response
 - Describe their statistics within each local window
 - Because texture is not entirely local. We need to see a few dots to describe it as dotted. Ditto for lined, checkered
 - But can't be too large, otherwise the description wouldn't change
 - The choice of scale is important for description

Texture representation: example



original image




derivative filter responses, squared

	<u>mean d/dx value</u>	<u>mean d/dy value</u>
<i>Win. #1</i>	4	10
<i>Win.#2</i>	18	7
⋮		
<i>Win.#9</i>	20	20
	⋮	

statistics to summarize patterns in small windows

Texture representation: example



	<u>mean d/ dx value</u>	<u>mean d/ dy value</u>
Win. #1	4	10
Win.#2	18	7
⋮		
Win.#9	20	20
	⋮	


**statistics to summarize
patterns in small
windows**

Texture representation: example

Windows with primarily horizontal edges

Both

Dimension 2 (mean d/dy value)



Dimension 1 (mean d/dx value)

Windows with small gradient in both directions

Windows with primarily vertical edges




	<u>mean d/dx value</u>	<u>mean d/dy value</u>
Win. #1	4	10
Win.#2	18	7
⋮		
Win.#9	20	20
	⋮	

statistics to summarize patterns in small windows


Texture representation: example



original image

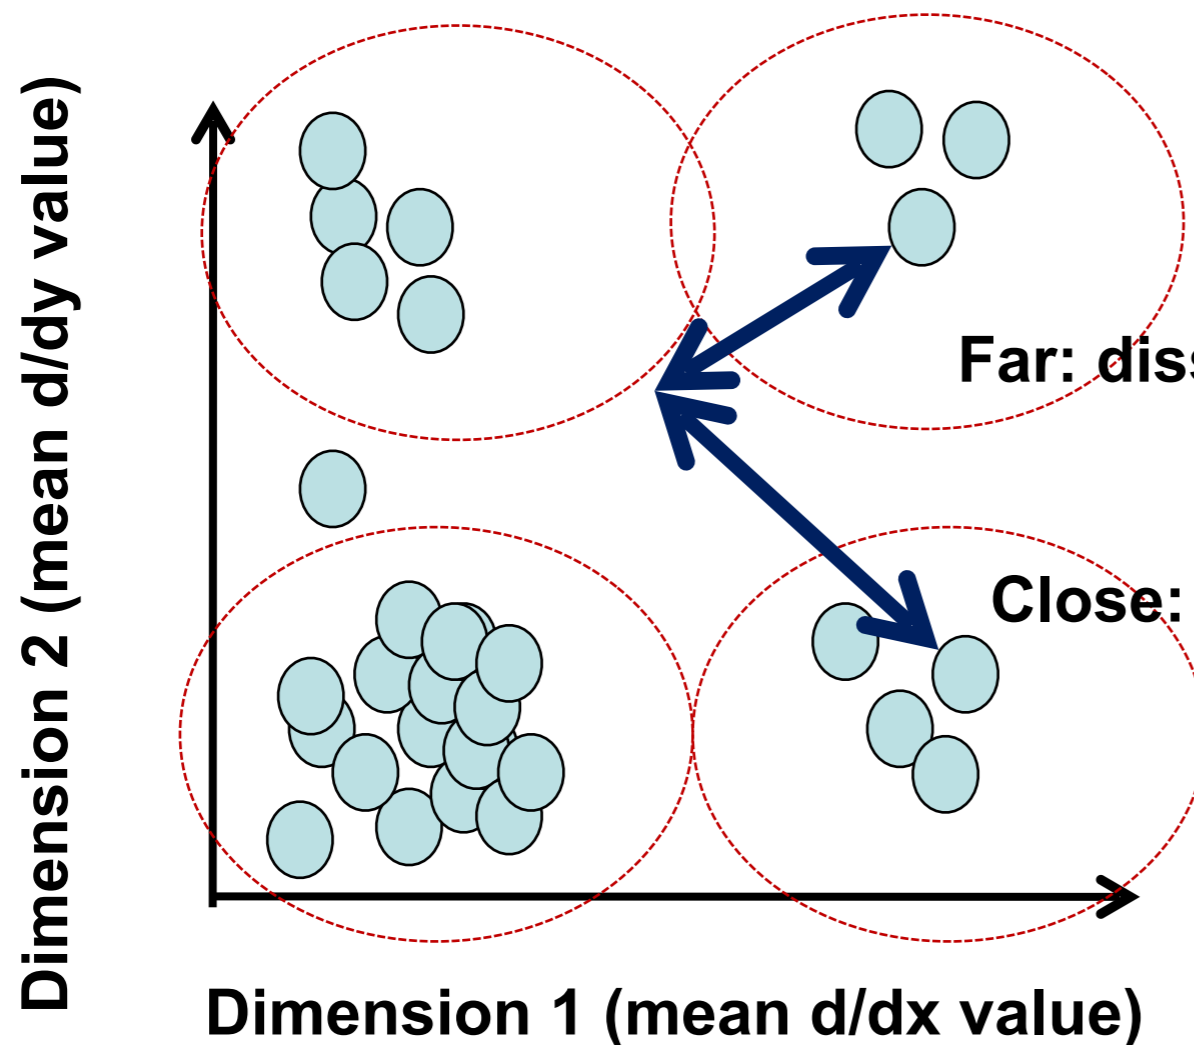


**derivative filter
responses, squared**



**visualization of the
assignment to texture
"types"**

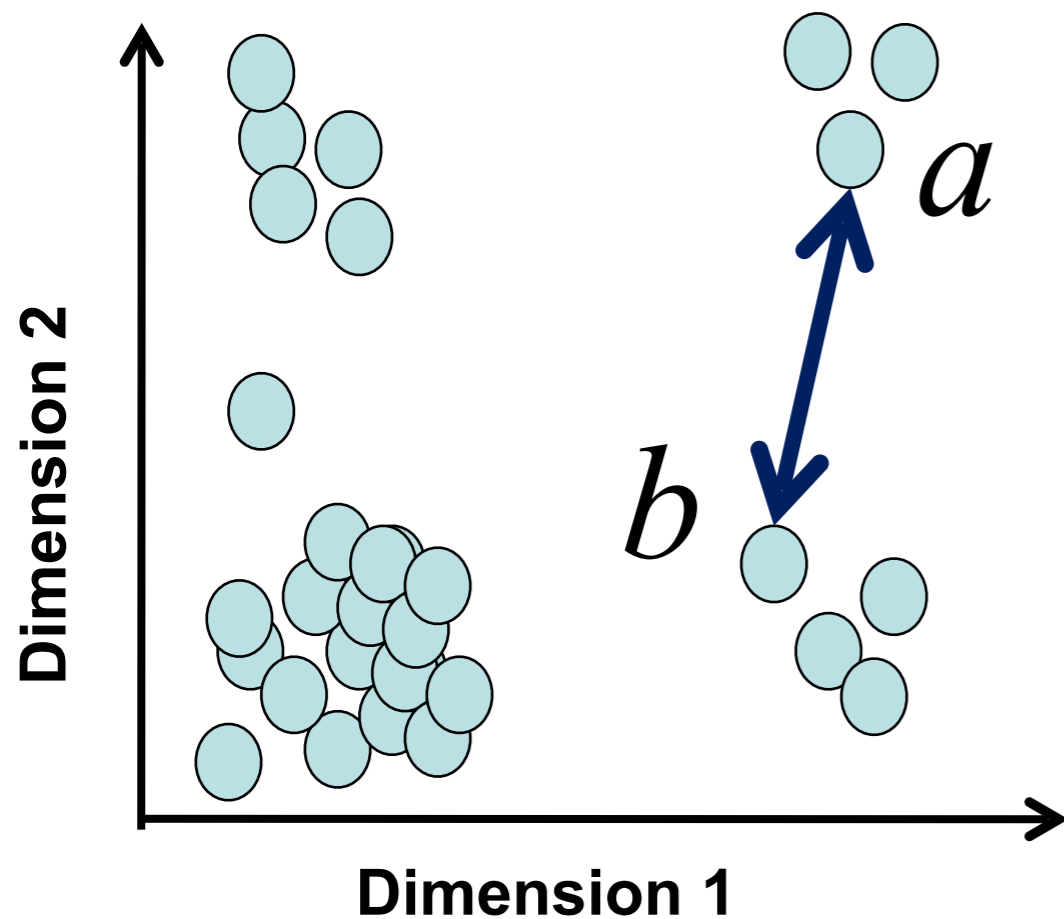
Texture representation: example



	<u>mean d/dx value</u>	<u>mean d/dy value</u>
<i>Win. #1</i>	4	10
<i>Win. #2</i>	18	7
⋮		
<i>Win. #9</i>	20	20
	⋮	

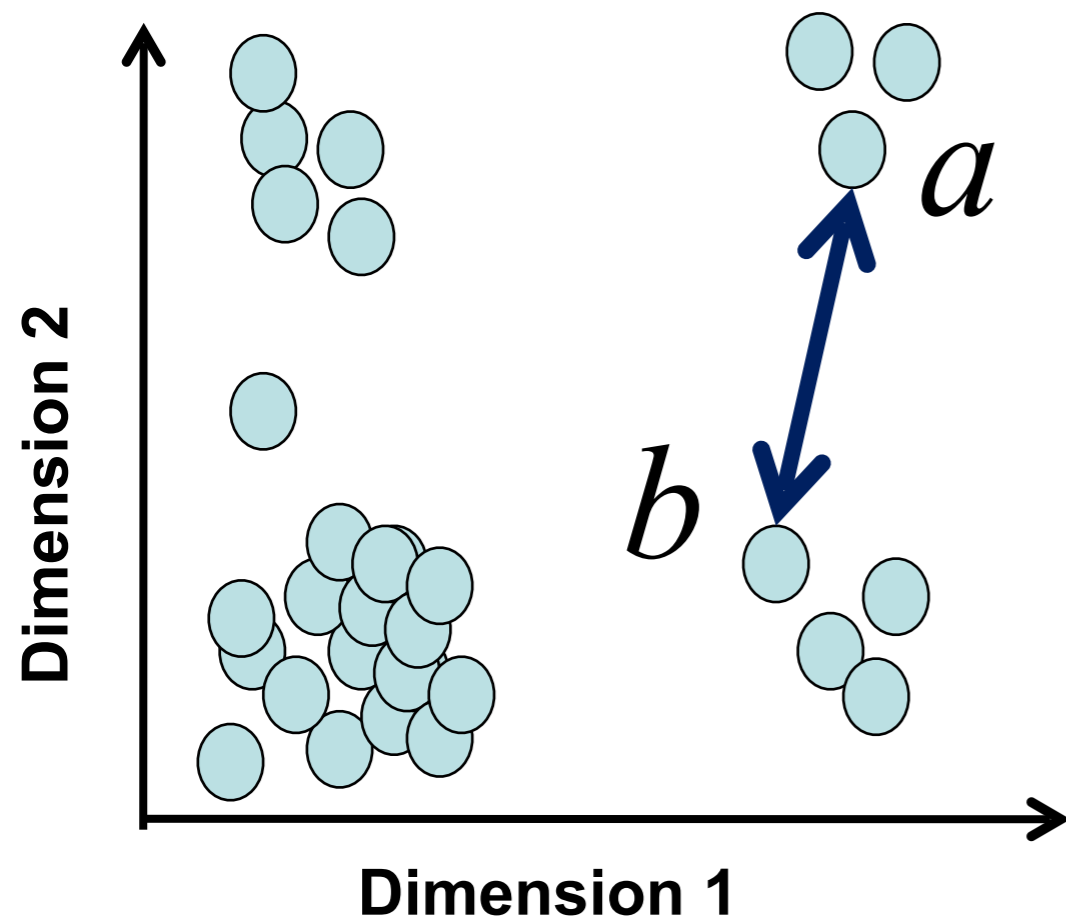
statistics to summarize
patterns in small
windows

Texture representation: example




$$D(a, b) = \sqrt{(a_1 - b_1)^2 + (a_2 - b_2)^2}$$

Texture representation: example




Distance reveals how dissimilar texture from window *a* is from texture in window *b*.



Texture representation: window scale

- We're assuming we know the relevant window size for which we collect these statistics.




Possible to perform scale selection by looking for window scale where texture description not changing.

Filter banks

- Our previous example used two filters, and resulted in a 2-dimensional feature vector to describe texture in a window.
 - x and y derivatives revealed something about local structure.
- We can generalize to apply a collection of multiple (d) filters: a “filter bank”
- Then our feature vectors will be d -dimensional.
 - still can think of nearness, farness in feature space

Filter banks

- What filters to put in the bank?
 - Typically we want a combination of scales and orientations, different types of patterns.




Matlab code available for these examples:

<http://www.robots.ox.ac.uk/~vgg/research/texclass/filters.html>

Multivariate Gaussian

$$p(x; \mu, \Sigma) = \frac{1}{(2\pi)^{n/2} |\Sigma|^{1/2}} \exp \left(-\frac{1}{2} (x - \mu)^T \Sigma^{-1} (x - \mu) \right).$$



$$\Sigma = \begin{bmatrix} 9 & 0 \\ 0 & 9 \end{bmatrix}$$

$$\Sigma = \begin{bmatrix} 16 & 0 \\ 0 & 9 \end{bmatrix}$$

$$\Sigma = \begin{bmatrix} 10 & 5 \\ 5 & 5 \end{bmatrix}$$

Filter bank

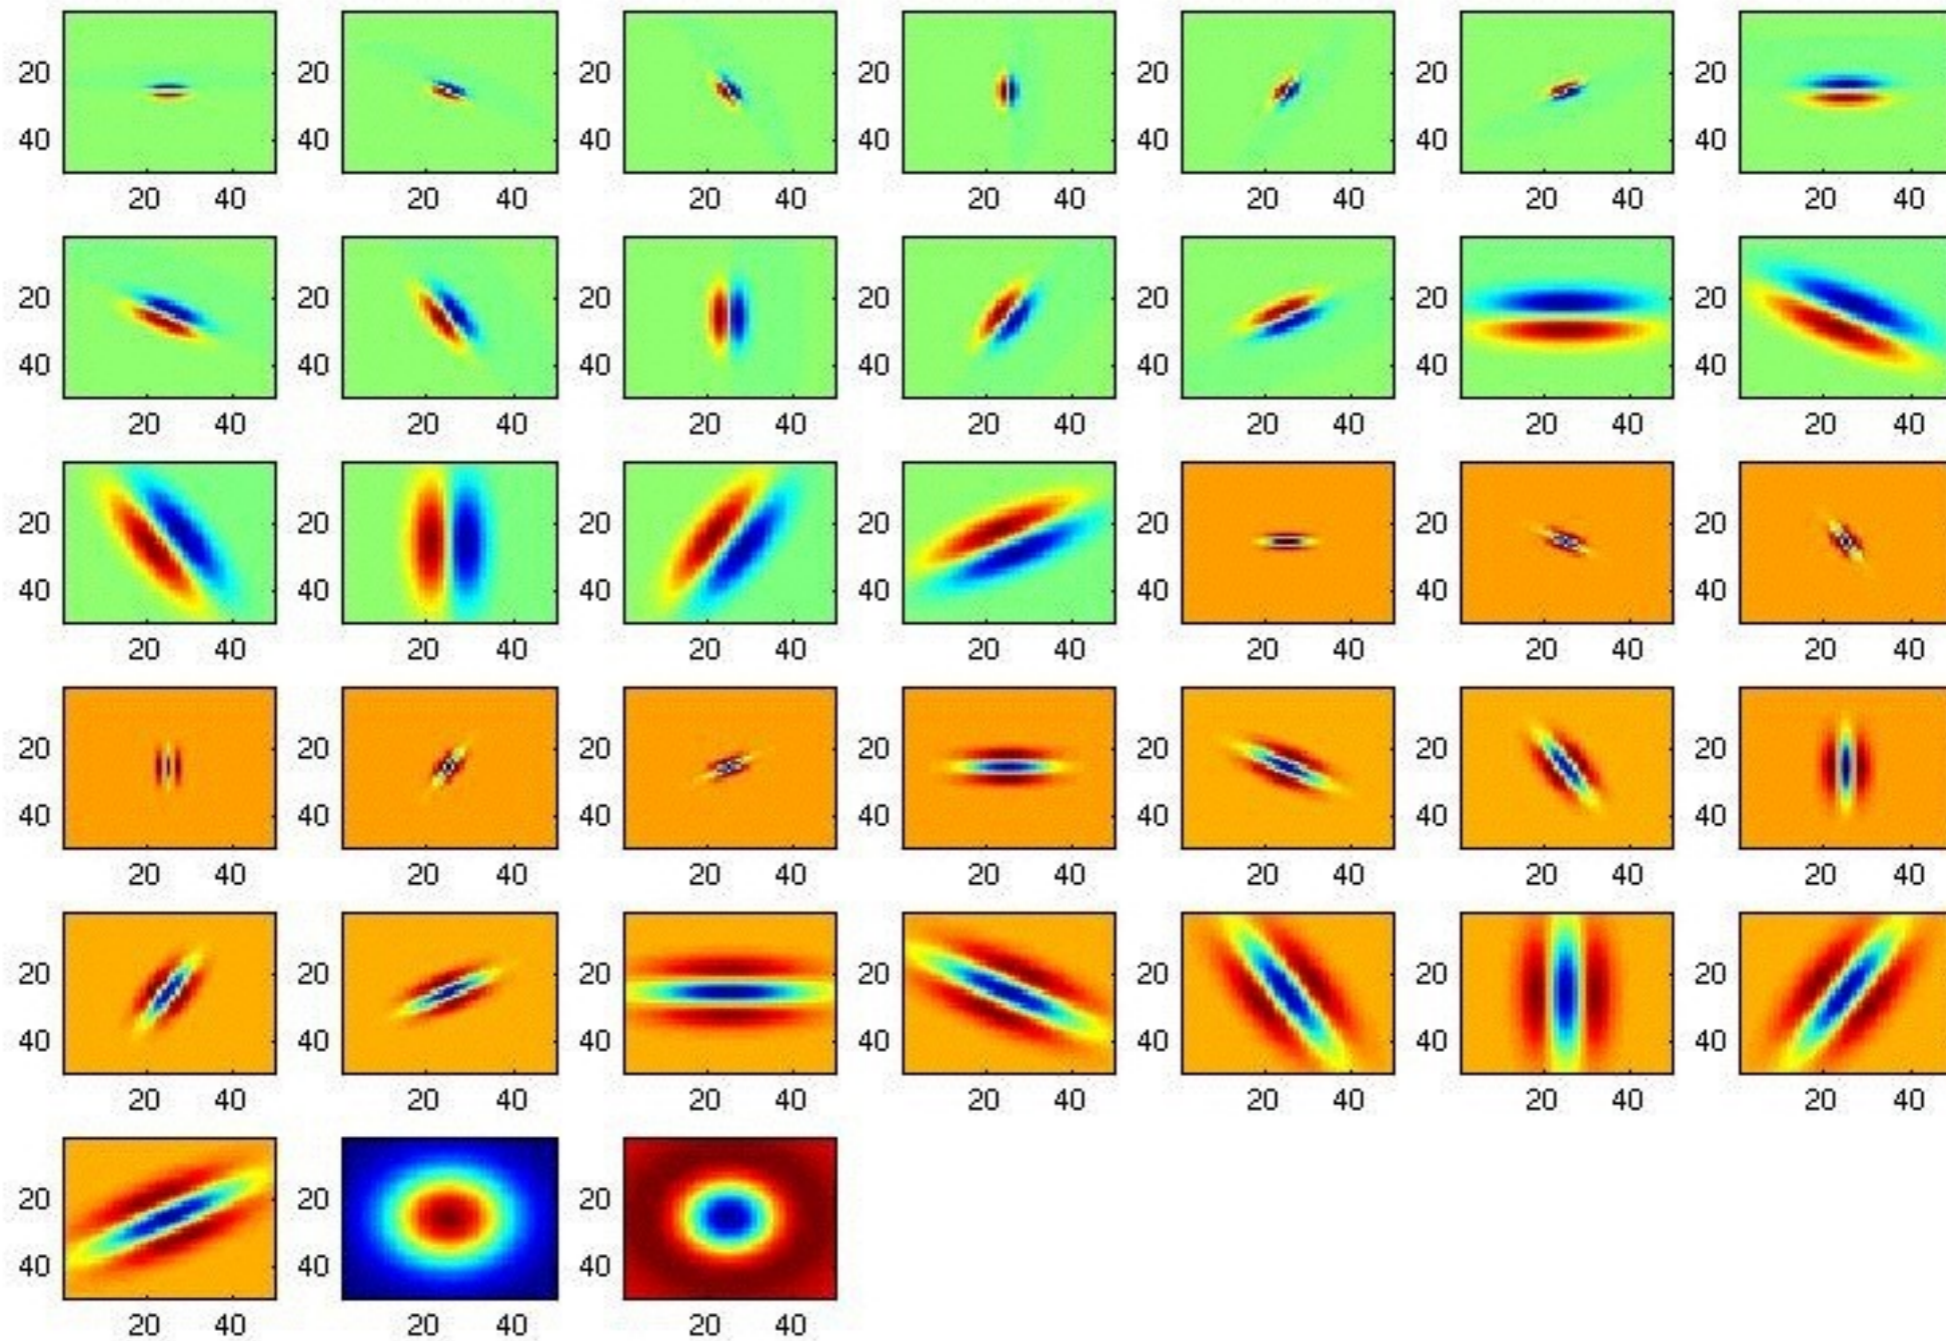








Image from <http://www.texasexplorer.com/austincap2.jpg>
































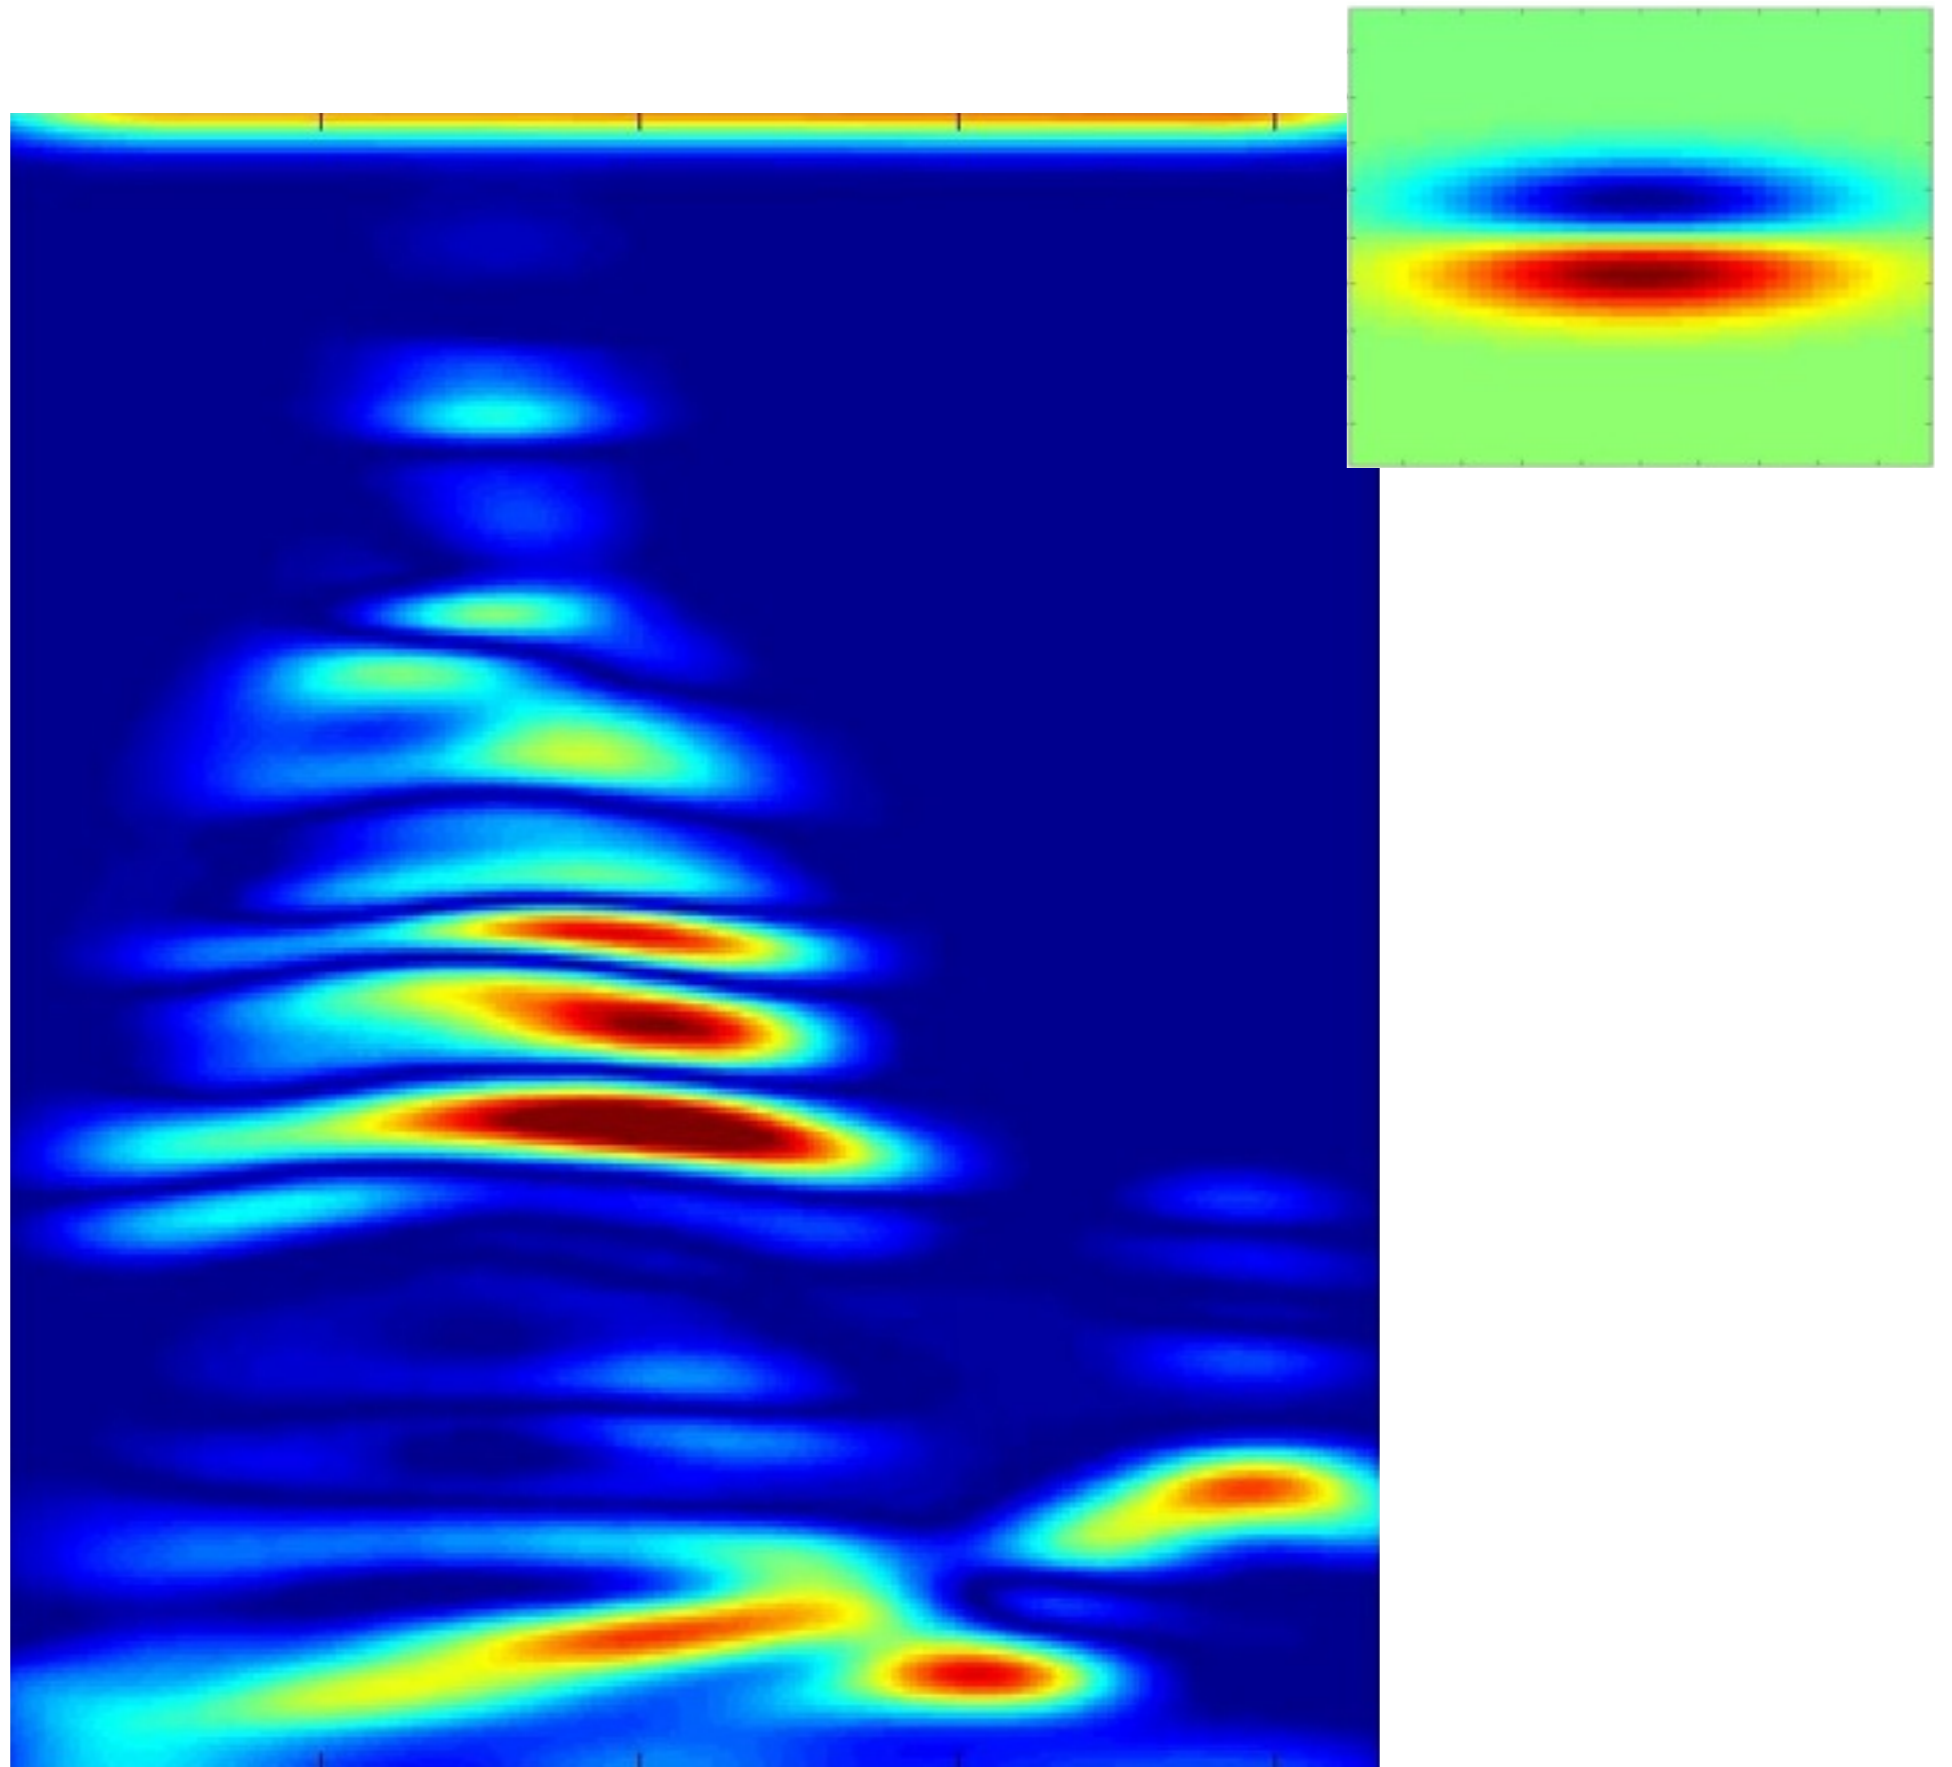








































[r1, r2, ..., r38]

We can form a feature vector from the list of responses at each pixel.

d -dimensional features

$$D(a, b) = \sqrt{\sum_{i=1}^d (a_i - b_i)^2}$$

Euclidean distance (L_2)



Counting in high dimensions

- Texture is a set of *textons* repeated in some way
- How do we find these repeated patterns?
- However, the representation is continuous so we cannot simply count the number of times we see a feature
- ***Vector quantization*** allows counting in high dimensions
 - Cluster the vectors into a fixed number of groups
 - Replace each vector with the the cluster center closest to it
- Often is better than binning.
- Each cluster is represented by a number, counting is easy.
- *Any* reasonable clustering method can be used.

K-means for vector quantization

Given a set of observations $(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$, where each observation is a d -dimensional real vector, k -means clustering aims to partition the \mathbf{n} observations into \mathbf{k} ($\leq \mathbf{n}$) sets $\mathbf{S} = \{S_1, S_2, \dots, S_k\}$ so as to minimize the within-cluster sum of squares (**WCSS**). In other words, its objective is to find:

$$\operatorname{argmin}_{\mathbf{S}} \sum_{i=1}^k \sum_{\mathbf{x} \in S_i} \|\mathbf{x} - \boldsymbol{\mu}_i\|^2$$

where $\boldsymbol{\mu}_i$ is the mean of points in S_i .


Easy to compute $\boldsymbol{\mu}$ given \mathbf{S} and vice versa.

Lloyd's algorithm for k-means

- Initialize k centers by picking k-points randomly
- Repeat till convergence (or max iterations)
 - Assign each point to the nearest center (assignment step)
 - Estimate the mean of each group (update step)

```
MATLAB    [idx, c] = kmeans(X, k)
```

K-means in action



Lloyd's algorithm for k-means

- Initialize k centers by picking k-points randomly
- Repeat till convergence (or max iterations)
 - Assign each point to the nearest center (assignment step)
 - Estimate the mean of each group (update step)

```
MATLAB [idx, c] = kmeans(X, k)
```


- Simple, fast and works well in practice
- But can be unstable
 - Run multiple times and the best solution (one with the smallest WCSS)
 - Better initializations are possible (e.g. kmeans++)

Textons in images


image




clustering into k=64 centers




convolution with f.b.



cluster (k-means)



square



aggregate



Uses of texture in vision: analysis

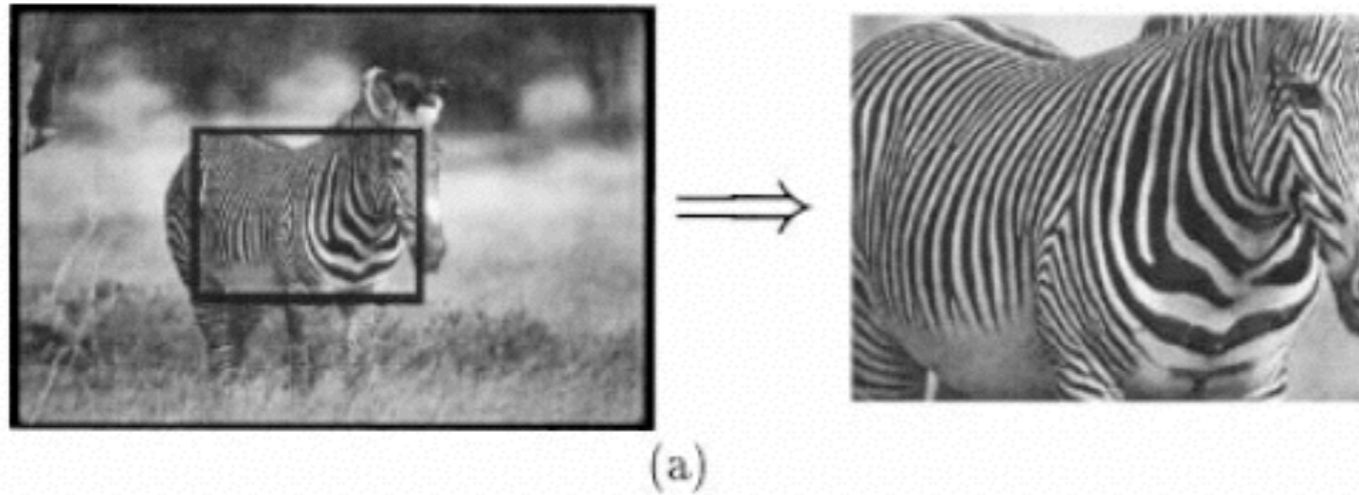
Classifying materials, “stuff”



Figure by Varma & Zisserman

Global texton histogram is a good representation

Texture features for image retrieval



1) 130066



2) 130070



3) 130068



4) 130051

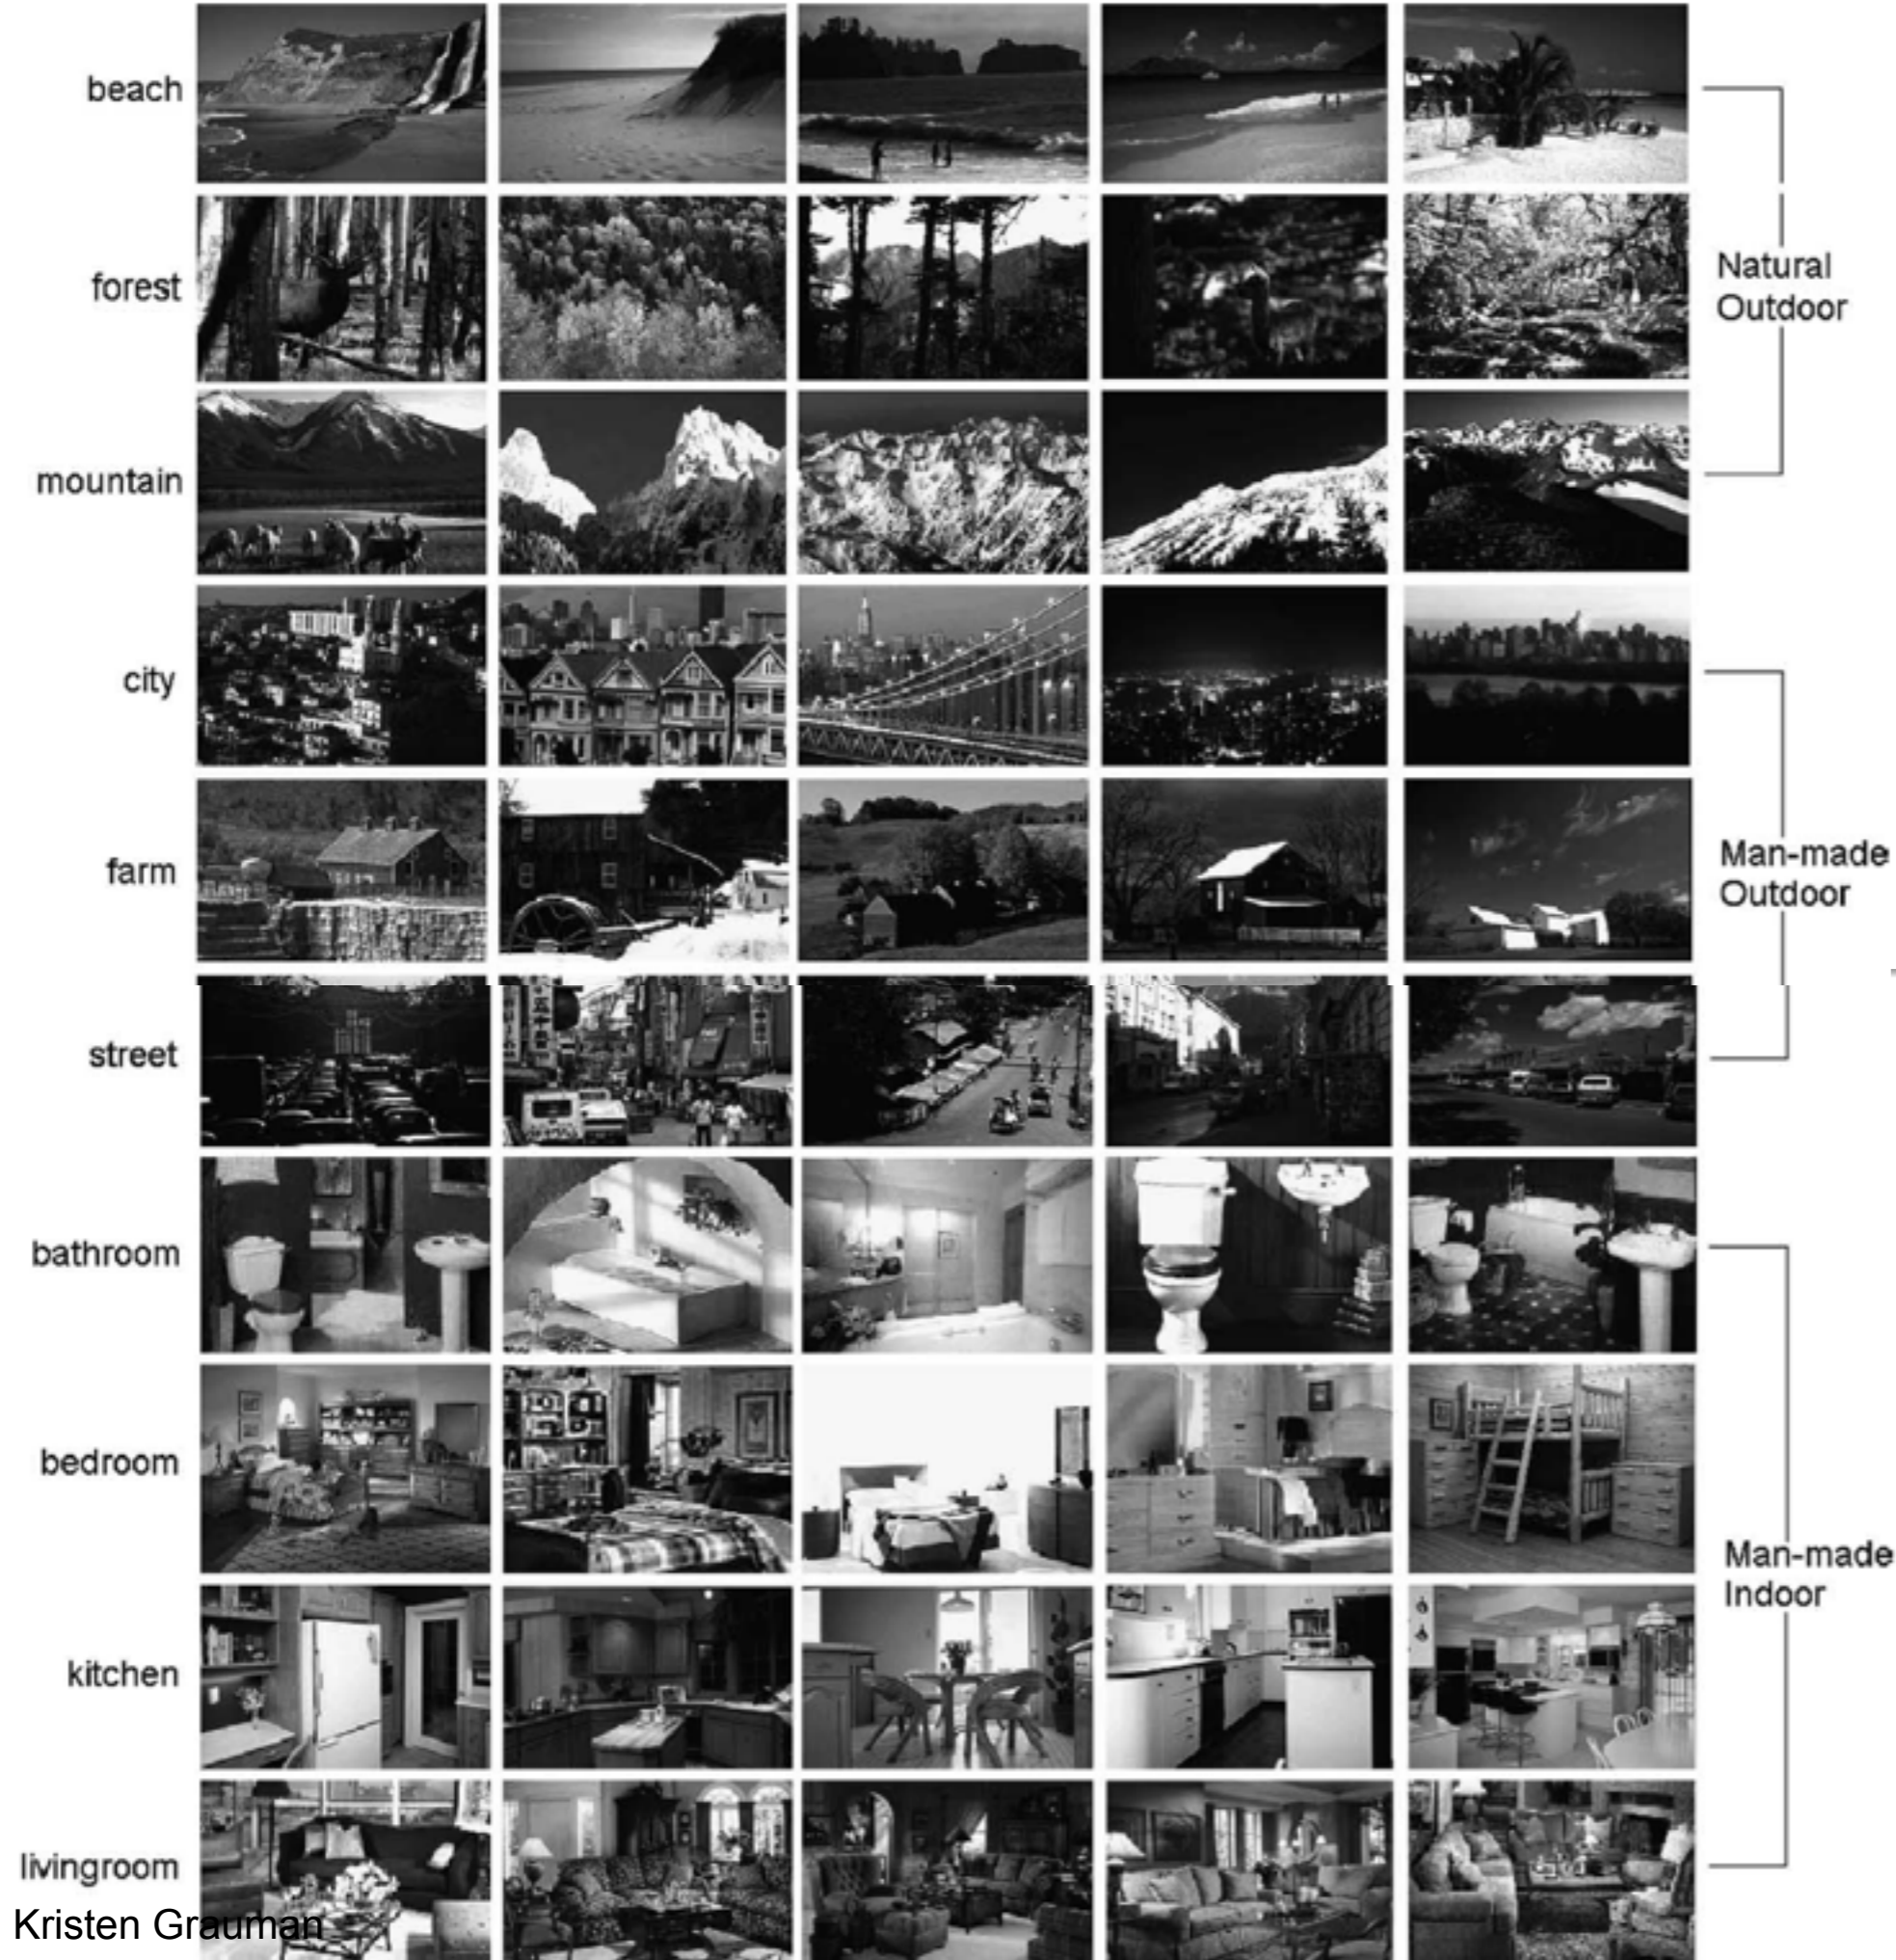


5) 130038



6) 130039

Y. Rubner, C. Tomasi, and L. J. Guibas. The earth mover's distance as a metric for image retrieval. *International Journal of Computer Vision*, 40(2):99-121, November 2000,



Characterizing scene categories by texture

L. W. Renninger and J. Malik. When is scene identification just texture recognition? *Vision Research* 44 (2004) 2301–2311



Segmenting
aerial imagery by
textures